

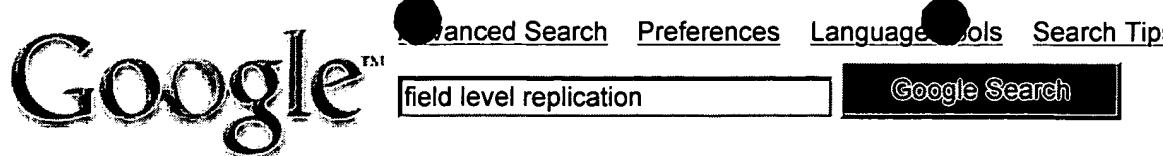
L Number	Hits	Search Text	DB	Time stamp
	1	"20020099728"	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/05 14: 25
	53	multi-valued adj attribute	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/05 11: 22
	9	(multi-valued adj attribute) and version and timestamp	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/05 11: 06
	1	multi-valued adj object adj attribute	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/05 11: 23
	0	multivalued adj object adj attribute	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/05 11: 27
	14	replication adj conflict	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/05 11: 37
	1	"5926816".PN.	USPAT; US-PGPUB	2003/05/05 11: 32
	1	"5884325".PN.	USPAT; US-PGPUB	2003/05/05 11: 33
	1	"5870765".PN.	USPAT; US-PGPUB	2003/05/05 11: 33
	1	"5806075".PN.	USPAT; US-PGPUB	2003/05/05 11: 34
	1	"5737601".PN.	USPAT; US-PGPUB	2003/05/05 11: 34
	52	(directory adj service) and (version adj number) and timestamp	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/05 12: 49
	1	"6266669".PN.	USPAT; US-PGPUB	2003/05/05 11: 41
	1	"5991771".PN.	USPAT; US-PGPUB	2003/05/05 11: 41
	1	"5924096".PN.	USPAT; US-PGPUB	2003/05/05 11: 42
	1	"5870733".PN.	USPAT; US-PGPUB	2003/05/05 11: 42
	1	"5867688".PN.	USPAT; US-PGPUB	2003/05/05 11: 42
	1	"5829001".PN.	USPAT; US-PGPUB	2003/05/05 11: 42
	1	"5796999".PN.	USPAT; US-PGPUB	2003/05/05 11: 44
	1	"5761500".PN.	USPAT; US-PGPUB	2003/05/05 11: 44
	1	"5758150".PN.	USPAT; US-PGPUB	2003/05/05 11: 44
	1	"5761500".PN.	USPAT; US-PGPUB	2003/05/05 11: 45
	1	"5758355".PN.	USPAT; US-PGPUB	2003/05/05 11: 45
	1	"5664228".PN.	USPAT; US-PGPUB	2003/05/05 11: 45
	1	"5666530".PN.	USPAT; US-PGPUB	2003/05/05 11: 45
	23	deletion adj indicator	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/05 12: 49
	805	707/203.ccis.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/05 14: 25

	4253	timestamp	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/05 14: 25
	88	707/203.ccls. and timestamp	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/05 14: 25
	15189	timestamp or (time adj stamp)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/05 14: 26
	185	707/203.ccls. and (timestamp or (time adj stamp))	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/05 14: 26
	57	(707/203.ccls. and (timestamp or (time adj stamp))) and @rlad<=20000631	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/05 14: 26
	163	replica adj object	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/05 16: 16
	54	(replica adj object) and version	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/06 09: 01
	407	replica\$4 adj object	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/06 09: 02
	31	(replica\$4 adj object) and version and timestamp	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/06 09: 20
	1	"5625818".PN.	USPAT; US-PGPUB	2003/05/06 09: 10
	1	"5613079".PN.	USPAT; US-PGPUB	2003/05/06 09: 10
	1	"5434994".PN.	USPAT; US-PGPUB	2003/05/06 09: 11
	1	"5386559".PN.	USPAT; US-PGPUB	2003/05/06 09: 12
	1	"5295222".PN.	USPAT; US-PGPUB	2003/05/06 09: 14
	1	"5016204".PN.	USPAT; US-PGPUB	2003/05/06 09: 14
	14	replication adj conflict	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/06 09: 32
	0	stated adj based adj replica\$4	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/06 09: 32
	1	state adj based adj replica\$4	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/06 09: 33
	5	(state adj based) with replica\$4	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/06 09: 48
	0	(highest adj version) and (latest adj timestamp)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/06 09: 49
	2	(higher adj version) and (latest adj timestamp)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/06 09: 49
	1	"5649195".PN.	USPAT; US-PGPUB	2003/05/06 12: 56
	1	"5627961".PN.	USPAT; US-PGPUB	2003/05/06 12: 56

	1	"5613113".PN.	USPAT; US-PGPUB	2003/05/06 12: 57
	1	"5600834".PN.	USPAT; US-PGPUB	2003/05/06 12: 57
	1	"5440730".PN.	USPAT; US-PGPUB	2003/05/06 13: 00
	1	"5261069".PN.	USPAT; US-PGPUB	2003/05/06 13: 01
	1	"5220657".PN.	USPAT; US-PGPUB	2003/05/06 13: 01
	27	"5765171"	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/07 11: 44
	1	"6233623".PN.	USPAT; US-PGPUB	2003/05/07 10: 42
	1	"6098078".PN.	USPAT; US-PGPUB	2003/05/07 10: 43
	1	"6065017".PN.	USPAT; US-PGPUB	2003/05/07 10: 43
	1	"5832487".PN.	USPAT; US-PGPUB	2003/05/07 10: 44
	1	"5787247".PN.	USPAT; US-PGPUB	2003/05/07 10: 44
	1	"5596574".PN.	USPAT; US-PGPUB	2003/05/07 10: 44
	15222	(timestamp or (time adj stamp))	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/07 11: 47
	1	(creation adj ((timestamp or (time adj stamp)))) and (version adj number) and (update adj ((timestamp or (time adj stamp))))	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/07 11: 48
	14	conflict adj resolution adj data	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/07 13: 21
	1	09/753998	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/07 13: 21
	21669	replica	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/07 16: 51
	5254	replica and version	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/07 16: 51
	20	(replica and version) and (creation adj time)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 09: 21
	256	active adj directory	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 08: 27
	35	(active adj directory) and replica	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 08: 27
	97	(active adj directory) and replica\$5	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 08: 27
	83	((active adj directory) and replica\$5) and version	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 08: 27
	41	((active adj directory) and replica\$5) and version) and (timestamp or (time adj stamp) or time-stamp)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 08: 36
	1		USPAT; US-PGPUB	2003/05/08 09: 16

	1	USPAT; US-PGPUB	2003/05/08 09:18
	1	USPAT; US-PGPUB	2003/05/08 09:18
	1	USPAT; US-PGPUB	2003/05/08 09:18
	1	USPAT; US-PGPUB	2003/05/08 09:18
	3 "5787441"	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 09:24
	1 "6138124"	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 09:36
	16 field adj level adj replica\$5	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 09:46
	2 field-level adj replica\$5	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 09:46
	7 row adj level adj replica\$5	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 09:48
	6 column adj level adj replica\$5	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 09:49
	40 ((US-6453324-\$ or US-6192365-\$ or US-6052681-\$ or US-6532479-\$ or US-6529917-\$ or US-6058401-\$ or US-5737601-\$ or US-5926816-\$ or US-5884325-\$ or US-6516327-\$ or US-6484185-\$ or US-5829001-\$ or US-4853843-\$ or US-6446092-\$ or US-6321236-\$ or US-6098078-\$ or US-5937413-\$ or US-5832487-\$ or US-6466980-\$ or US-6425016-\$ or US-5956718-\$ or US-5832514-\$ or US-5787262-\$ or US-5806074-\$ or US-5692182-\$ or US-5613079-\$).did. or (US-5434994-\$ or US-5893116-\$ or US-6539381-\$ or US-6343299-\$ or US-6301589-\$ or US-6023708-\$ or US-5968131-\$ or US-5961590-\$ or US-5909689-\$ or US-6138124-\$).did. or (US-20020029227-\$ or US-20020083083-\$ or US-20020129175-\$ or US-20010048728-\$).did.	USPAT; US-PGPUB	2003/05/08 14:45
	1 ((US-6453324-\$ or US-6192365-\$ or US-6052681-\$ or US-6532479-\$ or US-6529917-\$ or US-6058401-\$ or US-5737601-\$ or US-5926816-\$ or US-5884325-\$ or US-6516327-\$ or US-6484185-\$ or US-5829001-\$ or US-4853843-\$ or US-6446092-\$ or US-6321236-\$ or US-6098078-\$ or US-5937413-\$ or US-5832487-\$ or US-6466980-\$ or US-6425016-\$ or US-5956718-\$ or US-5832514-\$ or US-5787262-\$ or US-5806074-\$ or US-5692182-\$ or US-5613079-\$).did. or (US-5434994-\$ or US-5893116-\$ or US-6539381-\$ or US-6343299-\$ or US-6301589-\$ or US-6023708-\$ or US-5968131-\$ or US-5961590-\$ or US-5909689-\$ or US-6138124-\$).did. or (US-20020029227-\$ or US-20020083083-\$ or US-20020129175-\$ or US-20010048728-\$).did.) and (deletion adj indicator)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 14:51
	1 ((US-6453324-\$ or US-6192365-\$ or US-6052681-\$ or US-6532479-\$ or US-6529917-\$ or US-6058401-\$ or US-5737601-\$ or US-5926816-\$ or US-5884325-\$ or US-6516327-\$ or US-6484185-\$ or US-5829001-\$ or US-4853843-\$ or US-6446092-\$ or US-6321236-\$ or US-6098078-\$ or US-5937413-\$ or US-5832487-\$ or US-6466980-\$ or US-6425016-\$ or US-5956718-\$ or US-5832514-\$ or US-5787262-\$ or US-5806074-\$ or US-5692182-\$ or US-5613079-\$).did. or (US-5434994-\$ or US-5893116-\$ or US-6539381-\$ or US-6343299-\$ or US-6301589-\$ or US-6023708-\$ or US-5968131-\$ or US-5961590-\$ or US-5909689-\$ or US-6138124-\$).did. or (US-20020029227-\$ or US-20020083083-\$ or US-20020129175-\$ or US-20010048728-\$).did.) and (delet\$3 adj indicator)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 14:51
	83 delet\$3 adj indicator	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 14:53
	1 (delet\$3 adj indicator) and (null adj identifier)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 14:52

-	11	delet\$3 adj timestamp	USPAT; US-PGPUB; EPO; JPO; IBM_TDB USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/05/08 15:35
-	1	6212517.pn.		2003/05/08 15:35



Web · Images · Groups · Directory · News ·

Searched the web for **field level replication**.

Results 1 - 10 of about 251,000. Search took 0.17 seconds.

### BYTE.com

... But neither Oracle nor Microsoft supports **field-level replication**, which is promised for Notes 4.0, the Lotus groupware platform that's now slated for ...

[www.byte.com/art/9510/sec4/art3.htm](http://www.byte.com/art/9510/sec4/art3.htm) - 19k - May 6, 2003 - [Cached](#) - [Similar pages](#)

### BYTE.com

... Lotus has optimized **replication** in Notes 4.0. The most visible change is **field-level replication**. ... **Field-level replication** also helps in conflict resolution. ...

[www.byte.com/art/9604/sec12/art5.htm](http://www.byte.com/art/9604/sec12/art5.htm) - 15k - [Cached](#) - [Similar pages](#)

[ [More results from www.byte.com](#) ]

### Replication: a design principle for field research. Schafer, ...

... It is argued that the use of **replication** as a ... A great deal of research is done in **field** settings in ... State-level or district-based researchers, for example, are ...

[ericae.net/pare/getvn.asp?v=7&n=15](http://ericae.net/pare/getvn.asp?v=7&n=15) - 29k - [Cached](#) - [Similar pages](#)

### Introducing Application-level Replication and Naming into today's ...

... latter is that server-initiated **replication** can be ... semantic information making an application-level approach with ... also be another interesting **field** of research ...

[decweb.ethz.ch/WWW5/www331/overview.htm](http://decweb.ethz.ch/WWW5/www331/overview.htm) - 48k - [Cached](#) - [Similar pages](#)

### ED458217 2001-06-00 Replication in Field Research. ERIC Digest.

... **Replication** in Field Research ... A great deal of research is done in **field** settings in ... State-level or district-based researchers, for example, are often interested ...

[www.ericfacility.net/databases/ERIC\\_Digests/ed458217.html](http://www.ericfacility.net/databases/ERIC_Digests/ed458217.html) - 29k - [Cached](#) - [Similar pages](#)

### Synchrologic Data Sync provides mobile database synchronization ...

... Data Sync provides scaleable data **replication** between a server database and mobile users. A built-in rules engine and **field level synchronization** lets you ...

[www.synchrologic.com/about/\\_about\\_imobile\\_data\\_synchronization\\_details.html](http://www.synchrologic.com/about/_about_imobile_data_synchronization_details.html) - 25k - [Cached](#) - [Similar pages](#)

### database replication

... At this point I don't have a way to know WHICH **field** changed between the two replicas so either I do ROW level **replication** or I check each **field**. ...

[archives.postgresql.org/pgsql-hackers/1999-12/msg00914.php](http://archives.postgresql.org/pgsql-hackers/1999-12/msg00914.php) - [Similar pages](#)

### Sandshot Software - DBSync

... The method of synchronization is fully customizable down to the **field level**. ... It is different from database **replication** in that it does not attempt to keep ...

[www.sandshot.net/DBSync/DBSyncfaq.asp](http://www.sandshot.net/DBSync/DBSyncfaq.asp) - 21k - [Cached](#) - [Similar pages](#)

### Grameen Bank Replication Program

... The Dialogue provides promising replicators with a unique opportunity to spend a considerable amount of time at the **field level** and live through the daily ...

[www.grameen-info.org/grameen/gtrust/dialogue.html](http://www.grameen-info.org/grameen/gtrust/dialogue.html) - 3k - [Cached](#) - [Similar pages](#)

### Garagworks Home \$ Fields: What are they? Is there a list? \$ ...

... This is the basis for the Notes **field-level replication** model, and it's

one of the most ingenious solutions I've ever seen. \$Signature ...

domino.garageworks.com/domino%5Cgwhome.nsf/vGarageTools/ 572F2679809C206F882566D4007E66ED?

OpenDocument - 5k - [Cached](#) - [Similar pages](#)

Gooooooooogle ►

Result Page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [Next](#)

[Search within results](#)

Dissatisfied with your search results? [Help us improve.](#)

[Google Home](#) - [Advertise with Us](#) - [Business Solutions](#) - [Services & Tools](#) - [Jobs, Press, & Help](#)

©2003 Google



## Search Results

Search Results for: [replica\* object]

Found **271** of **108,822** searched.

**Warning: Maximum result set of 200 exceeded. Consider refining.**

### Search within Results

  > Advanced Search

> Search Help/Tips

**Sort by:** Title Publication Publication Date Score

**Results 1 - 20 of 200 short listing**

**Prev Page** 1 2 3 4 5 6 7 8 9 10 **Next Page**

- 1** A distributed 3D graphics library 100%  
 Blair MacIntyre , Steven Feiner  
**Proceedings of the 25th annual conference on Computer graphics and interactive techniques** July 1998
- 2** Performance evaluation of the Orca shared-object system 99%  
 Henri E. Bal , Raoul Bhoedjang , Rutger Hofman , Ceriel Jacobs , Koen Langendoen , Tim Rühl , M. Frans Kaashoek  
**ACM Transactions on Computer Systems (TOCS)** February 1998  
 Volume 16 Issue 1  
 Orca is a portable, object-based distributed shared memory (DSM) system. This article studies and evaluates the design choices made in the Orca system and compares Orca with other DSMs. The article gives a quantitative analysis of Orca's coherence protocol (based on write-updates with function shipping), the totally ordered group communication protocol, the strategy for object placement, and the all-software, user-space architecture. Performance measurements for 10 parallel applications ill ...
- 3** Eliminating synchronization bottlenecks using adaptive replication 98%  
 Martin C. Rinard , Pedro C. Diniz  
**ACM Transactions on Programming Languages and Systems (TOPLAS)** May 2003  
 Volume 25 Issue 3  
 This article presents a new technique, adaptive replication, for automatically eliminating synchronization bottlenecks in multithreaded programs that perform atomic operations on objects. Synchronization bottlenecks occur when multiple threads attempt to concurrently update the same object. It is often possible to eliminate synchronization bottlenecks by replicating objects. Each thread can then

update its own local replica without synchronization and without interacting with other threads. When ...

## 4 System support for object groups





## Search Results

Search Results for: [timestamp<AND>((version<AND>((replica\* object) ))) ]  
 Found 49 of 108,822 searched.

### Search within Results

  > Advanced Search

> Search Help/Tips

Sort by: Title Publication Publication Date Score

**Results 1 - 20 of 49 short listing**

Prev Page 1 2 3   
 Next Page

- 1** Dynamic quorum adjustment for partitioned data 93%

Maurice Herlihy  
**ACM Transactions on Database Systems (TODS)** June 1987  
 Volume 12 Issue 2

A partition occurs when functioning sites in a distributed system are unable to communicate. This paper introduces a new method for managing replicated data objects in the presence of partitions. Each operation provided by a replicated object has a set. of quorums, which are sets of sites whose cooperation suffices to execute the operation. The method permits an object's quorums to be adjusted dynamically in response to failures and recoveries. A transaction that is unable to progress using ...

- 2** A quorum-consensus replication method for abstract data types 90%

Maurice Herlihy  
**ACM Transactions on Computer Systems (TOCS)** February 1986  
 Volume 4 Issue 1

Replication can enhance the availability of data in distributed systems. This paper introduces a new method for managing replicated data. Unlike many methods that support replication only for uninterpreted files, this method systematically exploits type-specific properties of objects such as sets, queues, or directories to provide more effective replication. Each operation requires the cooperation of a certain number of sites for its successful completion. A quorum for an operation is any s ...

- 3** A shared, segmented memory system for an object-oriented database 88%

Mark F. Hornick , Stanley B. Zdonik  
**ACM Transactions on Information Systems (TOIS)** January 1987  
 Volume 5 Issue 1

This paper describes the basic data model of an object-oriented database and the basic architecture of the system implementing it. In particular, a secondary storage

segmentation scheme and a transaction-processing scheme are discussed. The segmentation scheme allows for arbitrary clustering of objects, including duplicates. The transaction scheme allows for many different sharing protocols ranging from those that enforce serializability to those that are nonserializable and require communi ...

- 4 Manageability, availability, and performance in porcupine: a highly scalable, cluster-based mail service** 85%

 Yasushi Saito , Brian N. Bershad , Henry M. Levy

**ACM Transactions on Computer Systems (TOCS) August 2000**

Volume 18 Issue 3

This paper describes the motivation, design and performance of Porcupine, a scalable mail server. The goal of Porcupine is to provide a highly available and scalable electronic mail service using a large cluster of commodity PCs. We designed Porcupine to be easy to manage by emphasizing dynamic load balancing, automatic configuration, and graceful degradation in the presence of failures. Key to the system's manageability, availability, and performance is that sessions, data, and underlying ...

- 5 Manageability, availability and performance in Porcupine: a highly scalable, cluster-based mail service** 83%

 Yasushi Saito , Brian N. Bershad , Henry M. Levy

**ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM symposium on Operating systems principles December 1999**

Volume 33 Issue 5

This paper describes the motivation, design, and performance of Porcupine, a scalable mail server. The goal of Porcupine is to provide a highly available and scalable electronic mail service using a large cluster of commodity PCs. We designed Porcupine to be easy to manage by emphasizing dynamic load balancing, automatic configuration, and graceful degradation in the presence of failures. Key to the system's manageability, availability, and performance is that sessions, data, and underlying serv ...

- 6 Concurrency versus availability: atomicity mechanisms for replicated data** 83%

 Maurice Herlihy

**ACM Transactions on Computer Systems (TOCS) August 1987**

Volume 5 Issue 3

A replicated object is a typed data object that is stored redundantly at multiple locations to enhance availability. Most techniques for managing replicated data have a two-level structure: At the higher level, a replica-control protocol reconstructs the object's state from its distributed components, and at the lower level, a standard concurrency-control protocol synchronizes accesses to the individual components. This paper explores an alternative approach to managing replicated data by p ...

- 7 An efficient and fault-tolerant solution for distributed mutual exclusion** 82%

 Divyakant Agrawal , Amr El Abbadi

**ACM Transactions on Computer Systems (TOCS) February 1991**

Volume 9 Issue 1

In this paper, we present an efficient and fault-tolerant algorithm for generating quorums to solve the distributed mutual exclusion problem. The algorithm uses a logical tree organization of the network to generate tree quorums, which are logarithmic in the size of the network in the best case. Our approach is resilient to both site and communication failures, even when such failures lead to network partitioning. Furthermore, the algorithm exhibits a property of graceful degradation,

i.e., ...

- 8** The generalized tree quorum protocol: an efficient approach for managing replicated data 82%

 D. Agrawal , A. El Abbadi

**ACM Transactions on Database Systems (TODS)** December 1992

Volume 17 Issue 4

In this paper, we present a low-cost fault-tolerant protocol for managing replicated data. We impose a logical tree structure on the set of copies of an object and develop a protocol that uses the information available in the logical structure to reduce the communication requirements for read and write operations. The tree quorum protocol is a generalization of the static voting protocol with two degrees of freedom for choosing quorums. In general, this results in significantly lower commun ...

- 9** Apologizing versus asking permission: optimistic concurrency control for abstract data types 80%

 M. Herlihy

**ACM Transactions on Database Systems (TODS)** March 1990

Volume 15 Issue 1

An optimistic concurrency control technique is one that allows transactions to execute without synchronization, relying on commit-time validation to ensure serializability. Several new optimistic concurrency control techniques for objects in decentralized distributed systems are described here, their correctness and optimality properties are proved, and the circumstances under which each is likely to be useful are characterized. Unlike many methods that classify operations only a ...

- 10** Position papers: Scalable information sharing in large scale distributed systems 80%

 Mustaque Ahamed , Sumeer Bhola , Rammohan Kordale , Francisco Torres-Rojas

**Proceedings of the seventh workshop on ACM SIGOPS European workshop: Systems support for worldwide applications** September 1996

Many application domains have already demonstrated that they can benefit greatly if efficient access can be provided to shared information across widely distributed users. We use the generic term *object* to describe units of shared information which could include files, web pages or language defined objects. Future applications will require object sharing modes richer than simple browsing. For example, a collaboration system that allows users distributed world-wide (e.g., managers of a mul ...

- 11** Concurrency and availability as dual properties of replicated atomic data 80%

 M. Herlihy

**Journal of the ACM (JACM)** April 1990

Volume 37 Issue 2

A replicated data object is a typed object that is stored redundantly at multiple locations in a distributed system. Each of the object's operations has a set of quorums, which are sets of sites whose cooperation is needed to execute that operation. A quorum assignment associates each operation with its set of quorums. An operation's quorums determine its availability, and the constraints governing an object's quorum assignments determine the range of availability properties realizable by r ...

- 12** Preserving and using context information in interprocess communication 80%

 Larry L. Peterson , Nick C. Buchholz , Richard D. Schlichting

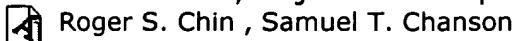
**ACM Transactions on Computer Systems (TOCS)** August 1989

## Volume 7 Issue 3

When processes in a network communicate, the messages they exchange define a partial ordering of externally visible events. While the significance of this partial order in distributed computing is well understood, it has not been made an explicit part of the communication substrate upon which distributed programs are implemented. This paper describes a new interprocess communication mechanism, called Psync, that explicitly encodes this partial ordering with each message. Th ...

**13** Distributed, object-based programming systems

80%



Roger S. Chin , Samuel T. Chanson

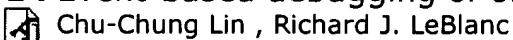
**ACM Computing Surveys (CSUR)** March 1991

Volume 23 Issue 1

The development of distributed operating systems and object-based programming languages makes possible an environment in which programs consisting of a set of interacting modules, or objects, may execute concurrently on a collection of loosely coupled processors. An object-based programming language encourages a methodology for designing and creating a program as a set of autonomous components, whereas a distributed operating system permits a collection of workstations or personal computers ...

**14** Event-based debugging of object/action programs

77%



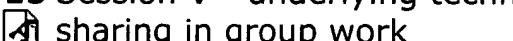
Chu-Chung Lin , Richard J. LeBlanc

**ACM SIGPLAN Notices , Proceedings of the 1988 ACM SIGPLAN and SIGOPS workshop on Parallel and distributed debugging** November 1988

Volume 24 Issue 1

**15** Session V - underlying technology for collaborative systems: Data sharing in group work

77%



Irene Greif , Sunil Sarin

**Proceedings of the 1986 ACM conference on Computer-supported cooperative work** December 1986

Data sharing is fundamental to computer-supported cooperative work: 'people share information through explicit communication channels and through their coordinated use of shared databases. Database support tools are therefore critical to the effective implementation of software for group work. This paper surveys data sharing requirements for group work, highlighting new database technologies that are especially likely to affect our ability to build computer systems supporting group work.

**16** Low cost management of replicated data in fault-tolerant distributed systems

77%



Thomas A. Joseph , Kenneth P. Birman

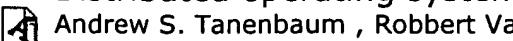
**ACM Transactions on Computer Systems (TOCS)** February 1986

Volume 4 Issue 1

Many distributed systems replicate data for fault tolerance or availability. In such systems, a logical update on a data item results in a physical update on a number of copies. The synchronization and communication required to keep the copies of replicated data consistent introduce a delay when operations are performed. In this paper, we describe a technique that relaxes the usual degree of synchronization, permitting replicated data items to be updated concurrently with other operations, ...

**17** Distributed operating systems

77%



Andrew S. Tanenbaum , Robbert Van Renesse

**ACM Computing Surveys (CSUR) December 1985**

Volume 17 Issue 4

Distributed operating systems have many aspects in common with centralized ones, but they also differ in certain ways. This paper is intended as an introduction to distributed operating systems, and especially to current university research about them. After a discussion of what constitutes a distributed operating system and how it is distinguished from a computer network, various key design issues are discussed. Then several examples of current research projects are examined in some detail ...

**18 Session 1: Nomadic computing and reconciliation: Application-** 77%

 **independent reconciliation for nomadic applications**

Marc Shapiro , Antony Rowstron , Anne-Marie Kermarrec

**Proceedings of the 9th workshop on ACM SIGOPS European workshop: beyond the PC: new challenges for the operating system September 2000**

We describe attempts to build an application-independent model to support reconciliation of diverged replicas of shared objects. While replicas are disconnected from one another, actions on the shared objects are recorded in a log. An action is composed of a precondition, an operation and a postcondition. When reconnecting, the system attempts to reconcile the divergent replicas, in several phases. A symbolic phase merges the separate logs, creating one or more schedules, such that preconditions ...

**19 Satellite-based information services: Cost based data dissemination in** 77%

 **satellite networks**

Bo Xu , Ouri Wolfson , Sam Chamberlain , Naphtali Rishé

**Mobile Networks and Applications January 2002**

Volume 7 Issue 1

We consider the problem of data dissemination in a broadcast network. In contrast to previously studied models, broadcasting is among peers, rather than client server. Such a model represents, for example, satellite communication among widely distributed nodes, sensor networks, and mobile ad hoc networks. We introduce a cost model for data dissemination in peer to peer broadcast networks. The model quantifies the tradeoff between the inconsistency of the data, and its transmission cost; the tran ...

**20 Disconnection modes for mobile databases** 77%

 **Joanne Holliday , Divyakant Agrawal , Amr El Abbadi**

**Wireless Networks July 2002**

Volume 8 Issue 4

As mobility permeates into todays computing and communication arena, we envision application infrastructures that will increasingly rely on mobile technologies. Traditional database applications and information service applications will need to integrate mobile entities: people and computers. In this paper, we develop a distributed database framework for mobile environments. A key requirement in such an environment is to support frequent connection and disconnection of database sites. We present ...

**Results 1 - 20 of 49****short listing**